EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	3408	707/202,203,204.ccls.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/07/13 13:19
L2	0	((network or server\$1 or host\$1 or master\$1) with (centra or ceter)) and (replicat\$4 or duplicat\$4 or cop\$3) and (conslidat\$3 or merg\$3 or combin\$1 combining) and ((freez\$3 or lock\$3 or prevent\$3) same (database\$1 or (data near base) or (data?base) or repositor\$3))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/07/13 13:24
L3	40	((network or server\$1 or host\$1 or master\$1) and (centra or ceter)) and (replicat\$4 or duplicat\$4 or cop\$3) and (conslidat\$3 or merg\$3 or combin\$1 combining) and ((freez\$3 or lock\$3 or prevent\$3) same (database\$1 or (data near base) or (data?base) or repositor\$3))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/07/13 14:05
L4	0	((network or server\$1 or host\$1 or master\$1) with(centra or ceter)) and (replicat\$4 or duplicat\$4 or cop\$3) and (conslidat\$3 or merg\$3 or combin\$1 combin\$5) and ((freez\$3 or lock\$3 or prevent\$3) same (database\$1 or (data near base) or (data?base) or repositor\$3))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/07/13 13:26
L5	651	leslie near wong	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/07/13 13:25
L6	16	5 and (replicat\$3 and freez\$)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/07/13 13:25

Page 1

EAST Search History

L7	0	((network or server\$1 or host\$1 or master\$1) with(centra or ceter)) and (replicat\$4 or duplicat\$4 or	US-PGPUB; USPAT; USOCR;	OR	ON	2006/07/13 13:27
		cop\$3) and (conslidat\$3 or merg\$3 or combin\$1 combin\$5) and ((freez\$3 or lock\$3 or prevent\$3 or quiesce) same (database\$1 or (data near base) or (data?base) or repositor\$3))	EPO; JPO; DERWENT; IBM_TDB			·
L8	1539	711/162.ccls.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/07/13 14:04
L9	1	(1 or 8) and (((network or server\$1 or host\$1 or master\$1) and (centra or ceter)) and (replicat\$4 or duplicat\$4 or cop\$3) and (conslidat\$3 or merg\$3 or combin\$1 combining) and ((freez\$3 or lock\$3 or prevent\$3) same (database\$1 or (data near base) or (data?base) or repositor\$3)))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/07/13 14:26
L10	1	"634993".apn.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/07/13 14:26
L11	0	"634993".apn. and (substantially)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/07/13 14:26
L12	0	"634993".apn. and (substant\$)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/07/13 14:27
L13	40488	((local network mirror) SAME (inhibit\$3 "slow down" reduc\$3) WITH (updat\$3 mirror\$3 replicat\$3 copy\$3))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/07/13 14:29

EAST Search History

L14	3931	((local network mirror) SAME (inhibit\$3 "slow down" reduc\$3) WITH (updat\$3 mirror\$3 replicat\$3 copy\$3) SAME (primary master center centra))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/07/13 14:30
L15	3931	((local network mirror) SAME ((inhibit\$3 "slow down" reduc\$3) WITH (updat\$3 mirror\$3 replicat\$3 copy\$3)) SAME (primary master center centra))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/07/13 14:31
L16	40	15 and (707/202,203,204.ccls. 711/162.ccls.)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/07/13 14:58
L17	83	((local network mirror) SAME ((inhibit\$3 "slow down" reduc\$3) WITH (updat\$3 mirror\$3 replicat\$3 copy\$3) WITH (react\$3 respond\$3 respons\$3)) SAME (primary master center centra))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/07/13 14:32
L18	0	(1 or 8) and (((local network mirror) SAME ((inhibit\$3 "slow down" reduc\$3) WITH (updat\$3 mirror\$3 replicat\$3 copy\$3) WITH (react\$3 respond\$3 respons\$3)) SAME (primary master center centra)))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/07/13 14:32
L19	0	(1 or 8) and (((local network mirror) SAME ((inhibit\$3 "slow down" reduc\$3) WITH (react\$3 respond\$3 respons\$3)) SAME (updat\$3 mirror\$3 replicat\$3 copy\$3) SAME (primary master center centra)))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/07/13 14:33



Subscribe (Full Service) Register (Limited Service, Free) Login

Search: • The ACM Digital Library • O The Guide

(local network mirror) same ((inhibit\$3 "slow down" slow redu





Feedback Report a problem Satisfaction survey

Terms used local network mirror same inhibit\$3 slow down slow reduc\$3 with updat\$3 mirror\$3 replicat\$3 copy\$3

Found 6,057 of 178,880

Sort results by

Display

results

relevance expanded form Save results to a Binder [?] Search Tips

Open results in a new

Try an Advanced Search Try this search in The ACM Guide

Results 1 - 20 of 200

window

Result page: 1 2 3 4 5 6 7 8 9 10

Relevance scale

Best 200 shown

1 Intrusion detection and modeling: Design space and analysis of worm defense

strategies

David Brumley, Li-Hao Liu, Pongsin Poosankam, Dawn Song

March 2006 Proceedings of the 2006 ACM Symposium on Information, computer and communications security ASIACCS '06

Publisher: ACM Press

Full text available: pdf(723.02 KB) Additional Information: full citation, abstract, references, index terms

We give the first systematic investigation of the design space of worm defense system strategies. We accomplish this by providing a taxonomy of defense strategies by abstracting away implementation-dependent and approach-specific details and concentrating on the fundamental properties of each defense category. Our taxonomy and analysis reveals the key parameters for each strategy that determine its effectiveness. We provide a theoretical foundation for understanding how these parameters interact ...

Keywords: antibody, blacklisting, defense strategy analysis, local containment, proactive protection, worm propagation, worm taxonomy, worms

2 Measuring thin-client performance using slow-motion benchmarking

Jason Nieh, S. Jae Yang, Naomi Novik

February 2003 ACM Transactions on Computer Systems (TOCS), Volume 21 Issue 1

Publisher: ACM Press

Full text available: pdf(871.62 KB)

Additional Information: full citation, abstract, references, citings, index terms

Modern thin-client systems are designed to provide the same graphical interfaces and applications available on traditional desktop computers while centralizing administration and allowing more efficient use of computing resources. Despite the rapidly increasing popularity of these client-server systems, there are few reliable analyses of their performance. Industry standard benchmark techniques commonly used for measuring desktop system performance are ill-suited for measuring the performance of ...

Keywords: Thin-client computing, client-server, measurement methodology, multimedia

3 Cluster I/O with River: making the fast case common

•	Remzi H. Arpaci-Dusseau, Eric Anderson, Noah Treuhaft, David E. Culler, Joseph M. Hellerstein, David Patterson, Kathy Yelick May 1999 Proceedings of the sixth workshop on I/O in parallel and distributed systems Publisher: ACM Press Full text available: pdf(1.20 MB) Additional Information: full citation, references, citings, index terms	
4	GPGPU: general purpose computation on graphics hardware David Luebke, Mark Harris, Jens Krüger, Tim Purcell, Naga Govindaraju, Ian Buck, Cliff Woolley, Aaron Lefohn August 2004 Proceedings of the conference on SIGGRAPH 2004 course notes GRAPH '04 Publisher: ACM Press Full text available: pdf(63.03 MB) Additional Information: full citation, abstract	
5	The graphics processor (GPU) on today's commodity video cards has evolved into an extremely powerful and flexible processor. The latest graphics architectures provide tremendous memory bandwidth and computational horsepower, with fully programmable vertex and pixel processing units that support vector operations up to full IEEE floating point precision. High level languages have emerged for graphics hardware, making this computational power accessible. Architecturally, GPUs are highly parallel s A low-bandwidth network file system Athicha Muthitacharoen, Benjie Chen, David Mazières	
~	October 2001 ACM SIGOPS Operating Systems Review , Proceedings of the eighteenth ACM symposium on Operating systems principles SOSP '01, Volume 35 Issue 5 Publisher: ACM Press Full text available: pdf(1.29 MB) Additional Information: full citation, abstract, references, citings, index terms	
	Users rarely consider running network file systems over slow or wide-area networks, as the performance would be unacceptable and the bandwidth consumption too high. Nonetheless, efficient remote file access would often be desirable over such networksparticularly when high latency makes remote login sessions unresponsive. Rather than run interactive programs such as editors remotely, users could run the programs locally and manipulate remote files through the file system. To do so, however, wo	
6 ③	Distributed file systems: concepts and examples Eliezer Levy, Abraham Silberschatz December 1990 ACM Computing Surveys (CSUR), Volume 22 Issue 4 Publisher: ACM Press	
	Full text available: pdf(5.33 MB) Additional Information: full citation, abstract, references, citings, index terms, review	
	The purpose of a distributed file system (DFS) is to allow users of physically distributed computers to share data and storage resources by using a common file system. A typical configuration for a DFS is a collection of workstations and mainframes connected by a local area network (LAN). A DFS is implemented as part of the operating system of each of the connected computers. This paper establishes a viewpoint that emphasizes the dispersed structure and decentralization of both data and con	
7 ③	Protocol architectures: A framework for scalable global IP-anycast (GIA) Dina Katabi, John Wroclawski April 2001 ACM SIGCOMM Computer Communication Review, Volume 31 Issue 2 supplement	

Publisher: ACM Press

Full text available: pdf(3.30 MB) Additional Information: full citation, abstract, references

This paper proposes GIA, a scalable architecture for global IP-anycast. Existing designs for providing IP-anycast must either globally distribute routes to individual anycast groups, or confine each anycast group to a pre-configured topological region. The first approach does not scale because of excessive growth in the routing tables, whereas the second one severely limits the utility of the service. Our design scales by dividing interdomain anycast routing into two components. The first compo ...

Keywords: anycast, architecture, internet, routing, scalable

8	Facial	modeling	and	animation

Jörg Haber, Demetri Terzopoulos

August 2004 Proceedings of the conference on SIGGRAPH 2004 course notes GRAPH

Publisher: ACM Press

Full text available: pdf(18.15 MB) Additional Information: full citation, abstract

In this course we present an overview of the concepts and current techniques in facial modeling and animation. We introduce this research area by its history and applications. As a necessary prerequisite for facial modeling, data acquisition is discussed in detail. We describe basic concepts of facial animation and present different approaches including parametric models, performance-, physics-, and learning-based methods. State-of-the-art techniques such as muscle-based facial animation, mass-s ...

9 Identity Boxing: A New Technique for Consistent Global Identity

Douglas Thain

November 2005 Proceedings of the 2005 ACM/IEEE conference on Supercomputing SC

'05

Publisher: IEEE Computer Society

Full text available: pdf(393.19 KB)

Additional Information: full citation, abstract

Today, users of the grid may easily authenticate themselves to computing resources around the world using a public key security infrastructure. However, users are forced to employ a patchwork of local identities, each assigned by a different local authority. This forces each grid system to provide a mapping from global to local identities, creating a significant administrative burden and inhibiting many possibilities of data sharing. To remedy this, we introduce the technique of identity boxing. ...

10 Optimistic replication

Yasushi Saito, Marc Shapiro

March 2005 ACM Computing Surveys (CSUR), Volume 37 Issue 1

Publisher: ACM Press

Full text available: pdf(656.72 KB) Additional Information: full citation, abstract, references, index terms

Data replication is a key technology in distributed systems that enables higher availability and performance. This article surveys optimistic replication algorithms. They allow replica contents to diverge in the short term to support concurrent work practices and tolerate failures in low-quality communication links. The importance of such techniques is increasing as collaboration through wide-area and mobile networks becomes popular. Optimistic replication deploys algorithms not seen in tradition ...

Keywords: Replication, disconnected operation, distributed systems, large scale systems, optimistic techniques

11	Distrbuted VEEs: HyperSpector: virtual distributed monitoring environments for	Г
_	to the state of th	
9	Kenichi Kourai, Shigeru Chiba	
	June 2005 Proceedings of the 1st ACM/USENIX international conference on Virtual	
	execution environments Publisher: ACM Press	
	Full text available: pdf(262.72 KB) Additional Information: full citation, abstract, references, index terms	
	In this paper, a virtual distributed monitoring environment called <i>HyperSpector</i> is described that achieves secure intrusion detection in distributed computer systems. While multiple intrusion detection systems (IDSes) can protect a distributed system from attackers, they can increase the number of insecure points in the protected system. HyperSpector overcomes this problem without any additional hardware by using virtualization to isolate each IDS from the servers it monitors. The IDSes a	
	Keywords: distributed IDS, inter-VM monitoring, virtual machine, virtual network	
12	Special issue on knowledge representation	
②	Developed Developed District Control	
	Publisher: ACM Press	
	Full text available: pdf(13.13 MB) Additional Information: full citation, abstract	
	In the fall of 1978 we decided to produce a special issue of the SIGART Newsletter devoted to a survey of current knowledge representation research. We felt that there were twe useful functions such an issue could serve. First, we hoped to elicit a clear picture of how people working in this subdiscipline understand knowledge representation research, to illuminate the issues on which current research is focused, and to catalogue what approaches and techniques are currently being developed. Secon	
13	Fast detection of communication patterns in distributed executions	
	Thomas Kunz, Michiel F. H. Seuren November 1997 Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research Publisher: IBM Press	
	Full text available: pdf(4.21 MB) Additional Information: full citation, abstract, references, index terms	
	Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun	
14	Mobile objects in distributed Oz	
\$	The state of the s	
	Publisher: ACM Press	
	Full text available: pdf(484.83 KB) Additional Information: full citation, abstract, references, citings, index terms	
	Some of the most difficult questions to answer when designing a distributed application are related to mobility: what information to transfer between sites and when and how to	

transfer it. Network-transparent distribution, the property that a program's behavior is independent of how it is partitioned among sites, does not directly address these questions. Therefore we propose to extend all language entities with a network behavior that enables efficient distributed programm ...

Keywords: latency tolerance, mobile objects, network transparency

15	Stateful distributed interposition John Reumann, Kang G. Shin February 2004 ACM Transactions on Computer Systems (TOCS). Volume 22 Issue 1	
	February 2004 ACM Transactions on Computer Systems (TOCS), Volume 22 Issue 1	
	Publisher: ACM Press Full text available: pdf(833.84 KB) Additional Information: full citation, abstract, references, index terms	
	Interposition-based system enhancements for multitiered servers are difficult to build because important system context is typically lost at application and machine boundaries. For example, resource quotas and user identities do not propagate easily between cooperating services that execute on different hosts or that communicate with each other via intermediary services. Application-transparent system enhancement is difficult to achieve when such context information is obscured by complex servic	
	Keywords : Distributed computing, component services, distributed context, multitiered services, operating systems, server consolidation	
16	A framework for scalable global IP-anycast (GIA)	
٨	Dina Katabi, John Wroclawski	
♥	August 2000 ACM SIGCOMM Computer Communication Review , Proceedings of the	
	conference on Applications, Technologies, Architectures, and Protocols for Computer Communication SIGCOMM '00, Volume 30 Issue 4	
	Publisher: ACM Press	
	Full text available: pdf(306.33 KB) Additional Information: full citation, abstract, references, citings, index terms	
	This paper proposes GIA, a scalable architecture for global IP-anycast. Existing designs for providing IP-anycast must either globally distribute routes to individual anycast groups, or confine each anycast group to a pre-configured topological region. The first approach does not scale because of excessive growth in the routing tables, whereas the second one severely limits the utility of the service. Our design scales by dividing interdomain anycast routing into two components. The first	
	Keywords: Internet, anycast, architecture, routing, scalable	
_	Level II technical support in a distributed computing environment Tim Leehane	
③	September 1996 Proceedings of the 24th annual ACM SIGUCCS conference on User services	
	Publisher: ACM Press Full text available: pdf(5.73 MB) Additional Information: full citation, references, index terms	
18	A High Availability Clustering Solution Phil Lewis August 1999 Linux Journal	

Publisher: Specialized Systems Consultants, Inc. Full text available: 1 html(34.77 KB) Additional Information: full citation, abstract, index terms Mr. Lewis tells us how he designed and implemented a simple high-availability solution for his company 19 ReVive: cost-effective architectural support for rollback recovery in shared-memory <u>multiprocessors</u> Milos Prvulovic, Zheng Zhang, Josep Torrellas May 2002 ACM SIGARCH Computer Architecture News, Proceedings of the 29th annual international symposium on Computer architecture ISCA '02, Proceedings of the 29th annual international symposium on Computer architecture ISCA '02, Volume 30 Issue 2 Publisher: IEEE Computer Society, ACM Press Full text available: pdf(1.38 MB) Additional Information: full citation, abstract, references, citings, index This paper presents ReVive, a novel general-purpose rollback recovery mechanism for shared-memory multiprocessors. ReVive carefully balances the conflicting requirements of availability, performance, and hardware cost. ReVive performs checkpointing, logging, and distributed parity protection, all memory-based. It enables recovery from a wide class of errors, including the permanent loss of an entire node. To maintain high performance, ReVive includes specialized hardware that performs frequent o ... **Keywords**: fault tolerance, shared-memory multiprocessors, rollback recovery, recovery, BER, logging, parity, checkpointing, availability 20 <u>Distance, dependencies, and delay in a global collaboration</u> James D. Herbsleb, Audris Mockus, Thomas A. Finholt, Rebecca E. Grinter December 2000 Proceedings of the 2000 ACM conference on Computer supported cooperative work **Publisher: ACM Press** Additional Information: full citation, abstract, references, citings, index Full text available: pdf(149.61 KB) Collaborations over distance must contend with the loss of the rich, subtle interactions that co-located teams use to coordinate their work. Previous research has suggested that one consequence of this loss is that cross-site work will take longer than comparable single-site work. We use both survey data and data from the change management system to measure the extent of delay in a multi-site software development organization. We also measure site interdependence, differences in same-site a ... Keywords: awareness, delay, global collaboration, informal communication, interdependence, software development, speed Results 1 - 20 of 200 Result page: 1 2 3 4 5 6 7 8 9 10 The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player



Subscribe (Full Service) Register (Limited Service, Free) Login

Search: • The ACM Digital Library • C The Guide

+"replicating" and database and ((local network mirror) same

SEARCH



Feedback Report a problem Satisfaction survey

Terms used

replicating and database and local network mirror same inhibit\$3 slow down slow reduc\$3 with updat\$3 mirror\$3 replicat\$3 copy\$3

 \triangle

Found 3,920 of 178,880

Sort results by

Display

results

relevance

expanded form

Save results to a Binder [?] Search Tips

Try an Advanced Search Try this search in The ACM Guide

Open results in a new window

Result page: 1 2 3 4 5 6 7 8 9 10

Results 1 - 20 of 200 Relevance scale Best 200 shown

1 Optimistic replication

Yasushi Saito, Marc Shapiro

March 2005 ACM Computing Surveys (CSUR), Volume 37 Issue 1

Publisher: ACM Press

Full text available: pdf(656.72 KB) Additional Information: full citation, abstract, references, index terms

Data replication is a key technology in distributed systems that enables higher availability and performance. This article surveys optimistic replication algorithms. They allow replica contents to diverge in the short term to support concurrent work practices and tolerate failures in low-quality communication links. The importance of such techniques is increasing as collaboration through wide-area and mobile networks becomes popular. Optimistic replication deploys algorithms not seen in tradition ...

Keywords: Replication, disconnected operation, distributed systems, large scale systems, optimistic techniques

2 Distributed file systems: concepts and examples

Eliezer Levy, Abraham Silberschatz

December 1990 ACM Computing Surveys (CSUR), Volume 22 Issue 4

Publisher: ACM Press

Full text available: pdf(5.33 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

The purpose of a distributed file system (DFS) is to allow users of physically distributed computers to share data and storage resources by using a common file system. A typical configuration for a DFS is a collection of workstations and mainframes connected by a local area network (LAN). A DFS is implemented as part of the operating system of each of the connected computers. This paper establishes a viewpoint that emphasizes the dispersed structure and decentralization of both data and con ...

3 Fast detection of communication patterns in distributed executions

Thomas Kunz, Michiel F. H. Seuren

November 1997 Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research

Publisher: IBM Press

Full text available: pdf(4.21 MB) Additional Information: full citation, abstract, references, index terms

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

4 GPGPU: general purpose computation on graphics hardware

David Luebke, Mark Harris, Jens Krüger, Tim Purcell, Naga Govindaraju, Ian Buck, Cliff Woolley, Aaron Lefohn

August 2004 Proceedings of the conference on SIGGRAPH 2004 course notes GRAPH '04

Publisher: ACM Press

Full text available: Def(63.03 MB) Additional Information: full citation, abstract

The graphics processor (GPU) on today's commodity video cards has evolved into an extremely powerful and flexible processor. The latest graphics architectures provide tremendous memory bandwidth and computational horsepower, with fully programmable vertex and pixel processing units that support vector operations up to full IEEE floating point precision. High level languages have emerged for graphics hardware, making this computational power accessible. Architecturally, GPUs are highly parallel s ...

5 Cluster I/O with River: making the fast case common

Remzi H. Arpaci-Dusseau, Eric Anderson, Noah Treuhaft, David E. Culler, Joseph M. Hellerstein, David Patterson, Kathy Yelick

May 1999 Proceedings of the sixth workshop on I/O in parallel and distributed systems

Publisher: ACM Press

Full text available: pdf(1.20 MB) Additional Information: full citation, references, citings, index terms

A low-bandwidth network file system

Athicha Muthitacharoen, Benjie Chen, David Mazières

October 2001 ACM SIGOPS Operating Systems Review, Proceedings of the eighteenth ACM symposium on Operating systems principles SOSP '01, Volume 35 Issue

Publisher: ACM Press

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(1.29 MB) terms

Users rarely consider running network file systems over slow or wide-area networks, as the performance would be unacceptable and the bandwidth consumption too high. Nonetheless, efficient remote file access would often be desirable over such networks--particularly when high latency makes remote login sessions unresponsive. Rather than run interactive programs such as editors remotely, users could run the programs locally and manipulate remote files through the file system. To do so, however, wo ...

7 Understanding fault-tolerant distributed systems

Flavin Cristian

February 1991 Communications of the ACM, Volume 34 Issue 2

Publisher: ACM Press

Additional Information: full citation, references, citings, index terms, Full text available: pdf(6.17 MB) review

8 ②	Research sessions: consistency and availability: Relaxed currency and consistency: how to say "good enough" in SQL Hongfei Guo, Per-Åke Larson, Raghu Ramakrishnan, Jonathan Goldstein June 2004 Proceedings of the 2004 ACM SIGMOD international conference on Management of data
	Publisher: ACM Press
	Full text available: pdf(606.81 KB) Additional Information: full citation, abstract, references, citings
	Despite the widespread and growing use of asynchronous copies to improve scalability, performance and availability, this practice still lacks a firm semantic foundation. Applications are written with some understanding of which queries can use data that is not entirely current and which copies are "good enough"; however, there are neither explicit requirements nor guarantees. We propose to make this knowledge available to the DBMS through explicit currency and consistency (C&C) constraints in qu
9	Stateful distributed interposition
٨	John Reumann, Kang G. Shin
~	February 2004 ACM Transactions on Computer Systems (TOCS), Volume 22 Issue 1 Publisher: ACM Press
	Full text available: pdf(833.84 KB) Additional Information: full citation, abstract, references, index terms
	Interposition-based system enhancements for multitiered servers are difficult to build because important system context is typically lost at application and machine boundaries. For example, resource quotas and user identities do not propagate easily between cooperating services that execute on different hosts or that communicate with each other via intermediary services. Application-transparent system enhancement is difficult to achieve when such context information is obscured by complex servic
	Keywords : Distributed computing, component services, distributed context, multitiered services, operating systems, server consolidation
10	LH* _{RS} a highly-available scalable distributed data structure
_	Witold Litwin, Rim Moussa, Thomas Schwarz
~	September 2005 ACM Transactions on Database Systems (TODS), Volume 30 Issue 3
	Publisher: ACM Press
	Full text available: pdf(774.32 KB) Additional Information: full citation, abstract, references, index terms
	LH*RS is a high-availability scalable distributed data structure (SDDS). An LH*RS file is hash partitioned over the distributed RAM of a multicomputer, for example, a network of PCs, and supports the unavailability of any k ≥ 1 of its server nodes. The value of k transparently grows with the file to offset the reliability decline. Only the number of the storage nodes potentially limits the file growth. The high-availability management uses a novel
	Keywords : P2P, Scalable distributed data structure, grid computing, high-availability, linear hashing, physical database design
11	Research session: database architectures for new hardware: Parallel querying with
•	non-dedicated computers
	Vijayshankar Raman, Wei Han, Inderpal Narang
	August 2005 Proceedings of the 31st international conference on Very large data
	bases VLDB '05 Publisher: VLDB Endowment

Full text available: pdf(462.68 KB) Additional Information: full citation, abstract, references, index terms

We present DITN, a new method of parallel querying based on dynamic outsourcing of join processing tasks to non-dedicated, heterogeneous computers. In DITN, partitioning is not the means of parallelism. Data layout decisions are taken outside the scope of the DBMS, and handled within the storage software; query processors see a "Data In The Network" image. This allows gradual scaleout as the workload grows, by using nondedicated computers. A typical operator in a parallel query pla ...

12 Bimodal multicast



Kenneth P. Birman, Mark Hayden, Oznur Ozkasap, Zhen Xiao, Mihai Budiu, Yaron Minsky May 1999 ACM Transactions on Computer Systems (TOCS), Volume 17 Issue 2

Publisher: ACM Press

Full text available: pdf(302.06 KB)

Additional Information: full citation, abstract, references, citings, index

There are many methods for making a multicast protocol "reliable." At one end of the spectrum, a reliable multicast protocol might offer tomicity guarantees, such as all-ornothing delivery, delivery ordering, and perhaps additional properties such as virtually synchronous addressing. At the other are protocols that use local repair to overcome transient packet loss in the network, offering "best effort" reliability. Yet none of this prior work has treated stability ...

13 Level II technical support in a distributed computing environment



Tim Leehane

September 1996 Proceedings of the 24th annual ACM SIGUCCS conference on User services

Publisher: ACM Press

Full text available: pdf(5.73 MB)

Additional Information: full citation, references, index terms

14 Management of a remote backup copy for disaster recovery



Richard P. King, Nagui Halim, Hector Garcia-Molina, Christos A. Polyzois

May 1991 ACM Transactions on Database Systems (TODS), Volume 16 Issue 2

Publisher: ACM Press

Full text available: pdf(2.48 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

A remote backup database system tracks the state of a primary system, taking over transaction processing when disaster hits the primary site. The primary and backup sites are physically isolated so that failures at one site are unlikely to propogate to the other. For correctness, the execution schedule at the backup must be equivalent to that at the primary. When the primary and backup sites contain a single processor, it is easy to achieve this property. However, this is harder to do when ...

Keywords: database initialization, hot spare, hot standby, remote backup

15 Replication in the harp file system



Barbara Liskov, Sanjay Ghemawat, Robert Gruber, Paul Johnson, Liuba Shrira September 1991 ACM SIGOPS Operating Systems Review, Proceedings of the thirteenth ACM symposium on Operating systems principles SOSP **'91**, Volume 25 Issue 5

Publisher: ACM Press

Full text available: pdf(1.60 MB)

Additional Information: full citation, abstract, references, citings, index <u>terms</u>

This paper describes the design and implementation of the Harp file system. Harp is a replicated Unix file system accessible via the VFS interface. It provides highly available and reliable storage for files and guarantees that file operations are executed atomically in spite of concurrency and failures. It uses a novel variation of the primary copy replication technique that provides good performance because it allows us to trade disk accesses for network communication. Harp is intended to be u ...

16 Scalable and fault-tolerant support for variable bit-rate data in the exedra streaming





Stergios V. Anastasiadis, Kenneth C. Sevcik, Michael Stumm

November 2005 ACM Transactions on Storage (TOS), Volume 1 Issue 4

Publisher: ACM Press

Full text available: pdf(1.01 MB) Additional Information: full citation, abstract, references, index terms

We describe the design and implementation of the Exedra continuous media server, and experimentally evaluate alternative resource management policies using a prototype system that we built. Exedra has been designed to provide scalable and efficient support for variable bit-rate media streams whose compression efficiency leads to reduced storage space and bandwidth requirements in comparison to constant bit-rate streams of equivalent quality. We examine alternative disk striping policies, and qua ...

Keywords: Content distribution, multimedia compression

17 Disaster recovery techniques for database systems





Manhoi Choy, Hong Va Leong, Man Hon Wong November 2000 Communications of the ACM

Publisher: ACM Press

Full text available: pdf(412.04 KB) Additional Information: full citation, references, index terms

18 Parallel multisource view maintenance

Xin Zhang, Lingli Ding, Elke A. Rundensteiner

January 2004 The VLDB Journal — The International Journal on Very Large Data

Bases, Volume 13 Issue 1

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(382.15 KB) Additional Information: full citation, abstract, index terms

In a distributed environment, materialized views are used to integrate data from different information sources and then store them in some centralized location. In order to maintain such materialized views, maintenance queries need to be sent to information sources by the data warehouse management system. Due to the independence of the information sources and the data warehouse, concurrency issues are raised between the maintenance queries and the local update transactions at each information so ...

Keywords: Concurrent data updates, Data warehousing, Parallel view maintenance, Performance evaluation

19 ReVive: cost-effective architectural support for rollback recovery in shared-memory





multiprocessors

Milos Prvulovic, Zheng Zhang, Josep Torrellas

May 2002 ACM SIGARCH Computer Architecture News, Proceedings of the 29th annual international symposium on Computer architecture ISCA '02, Proceedings of the 29th annual international symposium on Computer architecture ISCA '02, Volume 30 Issue 2

Publisher: IEEE Computer Society, ACM Press

Full text available: pdf(1.38 MB) Additional Information: full citation, abstract, references, citings, index Publisher Site

This paper presents ReVive, a novel general-purpose rollback recovery mechanism for shared-memory multiprocessors. ReVive carefully balances the conflicting requirements of availability, performance, and hardware cost. ReVive performs checkpointing, logging, and distributed parity protection, all memory-based. It enables recovery from a wide class of errors, including the permanent loss of an entire node. To maintain high performance, ReVive includes specialized hardware that performs frequent o ...

Keywords: fault tolerance, shared-memory multiprocessors, rollback recovery, recovery, BER, logging, parity, checkpointing, availability

20 Pen computing: a technology overview and a vision

André Meyer

July 1995 ACM SIGCHI Bulletin, Volume 27 Issue 3

Publisher: ACM Press

Additional Information: full citation, abstract, citings, index terms Full text available: pdf(5.14 MB)

This work gives an overview of a new technology that is attracting growing interest in public as well as in the computer industry itself. The visible difference from other technologies is in the use of a pen or pencil as the primary means of interaction between a user and a machine, picking up the familiar pen and paper interface metaphor. From this follows a set of consequences that will be analyzed and put into context with other emerging technologies and visions. Starting with a short historic ...

Results 1 - 20 of 200 Result page: 1 2 3 4 5 6 7 8 9 10

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat Q QuickTime Windows Media Player Real Player



Subscribe (Full Service) Register (Limited Service, Free) Login

Search: • The ACM Digital Library • O The Guide

(freez\$3 or split\$3) and (local mirror network) and (replicat\$3

SEARCH

HE ACTION OF ALLIERARY

Feedback Report a problem Satisfaction survey

Terms used

freez\$3 or split\$3 and local mirror network and replicat\$3 copy\$3 mirror\$3 updat\$3 and request

Found 4.481 of 178,880

Sort results

relevance Display expanded form \triangle

Save results to a Binder Search Tips

Try an Advanced Search Try this search in The ACM Guide

Open results in a new results window

Result page: 1 2 3 4 5 6 7 8 9 10

Relevance scale

Results 1 - 20 of 200

Best 200 shown

1 GPGPU: general purpose computation on graphics hardware

David Luebke, Mark Harris, Jens Krüger, Tim Purcell, Naga Govindaraju, Ian Buck, Cliff Woolley, Aaron Lefohn

August 2004 Proceedings of the conference on SIGGRAPH 2004 course notes GRAPH

Publisher: ACM Press

Full text available: pdf(63.03 MB) Additional Information: full citation, abstract

The graphics processor (GPU) on today's commodity video cards has evolved into an extremely powerful and flexible processor. The latest graphics architectures provide tremendous memory bandwidth and computational horsepower, with fully programmable vertex and pixel processing units that support vector operations up to full IEEE floating point precision. High level languages have emerged for graphics hardware, making this computational power accessible. Architecturally, GPUs are highly parallel s ...

² A <u>low-bandwidth network file system</u>

Athicha Muthitacharoen, Benjie Chen, David Mazières

October 2001 ACM SIGOPS Operating Systems Review , Proceedings of the eighteenth ACM symposium on Operating systems principles SOSP '01, Volume 35 Issue

Publisher: ACM Press

Full text available: pdf(1.29 MB)

Additional Information: full citation, abstract, references, citings, index

Users rarely consider running network file systems over slow or wide-area networks, as the performance would be unacceptable and the bandwidth consumption too high. Nonetheless, efficient remote file access would often be desirable over such networks--particularly when high latency makes remote login sessions unresponsive. Rather than run interactive programs such as editors remotely, users could run the programs locally and manipulate remote files through the file system. To do so, however, wo ...

3 Three-dimensional object recognition

Paul J. Besl, Ramesh C. Jain

March 1985 ACM Computing Surveys (CSUR), Volume 17 Issue 1

Publisher: ACM Press

Full text available: 規 pdf(7.76 MB)

Additional Information: full citation, abstract, references, citings, index

1	eı	rm	S,	rev	/iew

A general-purpose computer vision system must be capable of recognizing threedimensional (3-D) objects. This paper proposes a precise definition of the 3-D object recognition problem, discusses basic concepts associated with this problem, and reviews the relevant literature. Because range images (or depth maps) are often used as sensor input instead of intensity images, techniques for obtaining, processing, and characterizing range data are also surveyed.

4	Special session on knowledge-based programs: Causation, action, and counterfactuals Judea Pearl March 1996 Proceedings of the 6th conference on Theoretical aspects of rationality and knowledge Publisher: Morgan Kaufmann Publishers Inc. Full text available: pdf(1.69 MB) Additional Information: full citation, abstract, references							
	The central aim of many empirical studies in the physical, behavioral, social, and biological sciences is the elucidation of cause-effect relationships among variables. It is through cause-effect relationships that we obtain a sense of a "deep understanding" of a given phenomenon, and it is through such relationships that we obtain a sense of being "in control," namely, that we are able to shape the course of events by deliberate actions or policies. It is for these two reasons, understanding							
5 �	Performing remote operations efficiently on a local computer network Alfred Z. Spector April 1982 Communications of the ACM, Volume 25 Issue 4 Publisher: ACM Press							
	Full text available: pdf(1.58 MB) Additional Information: full citation, abstract, references, citings, index terms							
	A communication model is described that can serve as a basis for a highly efficient communication subsystem for local networks. The model contains a taxonomy of communication instructions that can be implemented efficiently and can be a good basis for interprocessor communication. These communication instructions, called remote references, cause an operation to be performed by a remote process and, optionally, cause a value to be returned. This paper also presents implementation considerati							
	Keywords: communication models, efficient communication, transactions							
6 �	High-speed switch scheduling for local-area networks Thomas E. Anderson, Susan S. Owicki, James B. Saxe, Charles P. Thacker November 1993 ACM Transactions on Computer Systems (TOCS), Volume 11 Issue 4							
	Publisher: ACM Press Full text available: pdf(2.37 MB) Additional Information: full citation, abstract, references, citings, index terms							

Current technology trends make it possible to build communication networks that can support high-performance distributed computing. This paper describes issues in the design of a prototype switch for an arbitrary topology point-to-point network with link speeds of up to 1 Gbit/s. The switch deals in fixed-length ATM-style cells, which it can process at a rate of 37 million cells per second. It provides high bandwidth and low latency for datagram traffic. In addition, it supports real-time t ...

Keywords: ATM networks, iterative matching, statistical matching, switching scheduling

11 <u>Local networks</u> William Stallings

March 1984 ACM Computing Surveys (CSUR), Volume 16 Issue 1

Publisher: ACM Press

...

Full text available: pdf(3.01 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

The rapidly evolving field of local network technology has produced a steady stream of local network products in recent years. The IEEE 802 standards that are now taking shape, because of their complexity, do little to narrow the range of alternative technical approaches and at the same time encourage more vendors into the field. The purpose of this paper is to present a systematic, organized overview of the alternative architectures for and design approaches to local networks.

12 Enabling full service surrogates using the portable channel representation Micah Beck, Terry Moore, Leif Abrahamsson, Christophe Achouiantz, Patrick Johansson April 2001 Proceedings of the 10th international conference on World Wide Web Publisher: ACM Press Full text available: pdf(282.92 KB) Additional Information: full citation, references, index terms Keywords: content distribution, dynamic content, mirroring, portability, replication, surrogate, web server 13 Distributed servers architecture for networked video services S.-H. Gary Chan April 2001 IEEE/ACM Transactions on Networking (TON), Volume 9 Issue 2 Publisher: IEEE Press Additional Information: full citation, references, citings, index terms, Full text available: pdf(300.02 KB) review Keywords: architecture, caching schemes, distributed servers, network channels and local storage, tradeoff, unicast and multicast, video-on-command

14 Run-time adaptation in river

Remzi H. Arpaci-Dusseau

February 2003 ACM Transactions on Computer Systems (TOCS), Volume 21 Issue 1

Publisher: ACM Press

Full text available: pdf(849.04 KB) Additional Information: full citation, abstract, references, index terms

We present the design, implementation, and evaluation of run-time adaptation within the River dataflow programming environment. The goal of the River system is to provide adaptive mechanisms that allow database query-processing applications to cope with performance variations that are common in cluster platforms. We describe the system and its basic mechanisms, and carefully evaluate those mechanisms and their effectiveness. In our analysis, we answer four previously unanswered and important que ...

Keywords: Performance availability, clusters, parallel I/O, performance faults, robust performance, run-time adaptation

\$	DNS: On the responsiveness of DNS-based network control Jeffrey Pang, Aditya Akella, Anees Shaikh, Balachander Krishnamurthy, Srinivasan Seshan October 2004 Proceedings of the 4th ACM SIGCOMM conference on Internet measurement Publisher: ACM Press Full text available: pdf(255.96 KB) Additional Information: full citation, abstract, references, index terms	
	For the last few years, large Web content providers interested in improving their scalability and availability have increasingly turned to three techniques: mirroring, content distribution, and ISP multihoming. The Domain Name System (DNS) has gained a prominent role in the way each of these techniques directs client requests to achieve the goals of scalability and availability. The DNS is thought to offer the transparent and agile control necessary to react quickly to ISP link failures or ph	
	Keywords: DNS, network control, time-to-live	
16	May 2003 Proceedings of the 1st international conference on Mobile systems,	
	applications and services MobiSys '03 Publisher: ACM Press	
	Full text available: pdf(232.79 KB) Additional Information: full citation, abstract, references	
	MobileIP, the standard for Internet mobility, enables transparent mobility for a mobile node, but requires communication to take a multihop path through the node's Home Agent. Although a user with a multiple-interface mobile node may desire the ability to communicate locally, perhaps while disconnected from the Internet, MobileIP offers no such support. Contact Networking provides lightweight, localized network communication to a node with diverse network interfaces. The goal is to provide suppor	
17	Serverless network file systems	
٥	The same E. Andreas and Affair at D. Dalatin, Barrian M. Marie, David A. Dalatinana Duran C.	
	Publisher: ACM Press	
	Full text available: pdf(2.69 MB) Additional Information: full citation, abstract, references, citings, index terms	
	We propose a new paradigm for network file system design: serverless network file systems. While traditional network file systems rely on a central server machine, a serverless system utilizes workstations cooperating as peers to provide all file system services. Any machine in the system can store, cache, or control any block of data. Our approach uses this location independence, in combination with fast local area networks, to provide better performance and scalability th	
	Keywords : RAID, log cleaning, log structured, log-based striping, logging, redundant data storage, scalable performance	

18 Serverless network file systems

T. E. Anderson, M. D. Dahlin, J. M. Neefe, D. A. Patterson, D. S. Roselli, R. Y. Wang December 1995 ACM STCORS Constitution of the control of December 1995 ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth ACM symposium on Operating systems principles SOSP '95, Volume 29

Publisher: ACM Press

Full text available: pdf(2.48 MB) Additional Information: full citation, references, citings, index terms

19	Michael D. Dahlin, Clifford J. Ma Patterson May 1994 ACM SIGMETRICS F 1994 ACM SIGMETR	the policies for scalable network file systems ther, Randolph Y. Wang, Thomas E. Anderson, David A. Performance Evaluation Review, Proceedings of the RICS conference on Measurement and modeling of SIGMETRICS '94, Volume 22 Issue 1	
	Publisher: ACM Press	SIGMETRICS 94, volume 22 issue i	
	Full text available: pdf(1.42 MB)	Additional Information: <u>full citation</u> , <u>abstract</u> , <u>references</u> , <u>citings</u> , <u>index</u> <u>terms</u>	
	Current network file system	protocols rely heavily on a central server to coordinate file	

activity among client workstations. This central server can become a bottleneck that limits scalability for environments with large numbers of clients. In central server systems such as NFS and AFS, all client writes, cache misses, and coherence messages are handled by the server. To keep up with this workload, expensive server machines are needed, configured with high-performance CPUs, memory systems, ...

20 Petal: distributed virtual disks

Edward K. Lee, Chandramohan A. Thekkath

September 1996 ACM SIGPLAN Notices, ACM SIGOPS Operating Systems Review, Proceedings of the seventh international conference on Architectural support for programming languages and operating systems ASPLOS-**VII**, Volume 31, 30 Issue 9, 5

Publisher: ACM Press

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(1.10 MB)

The ideal storage system is globally accessible, always available, provides unlimited performance and capacity for a large number of clients, and requires no management. This paper describes the design, implementation, and performance of Petal, a system that attempts to approximate this ideal in practice through a novel combination of features. Petal consists of a collection of network-connected servers that cooperatively manage a pool of physical disks. To a Petal client, this collection appear ...

Results 1 - 20 of 200 Result page: **1** <u>2</u> <u>3</u> <u>4</u> <u>5</u> <u>6</u> <u>7</u> <u>8</u> <u>9</u> <u>10</u>

The ACM Portal is published by the Association for Computing Machinery. Copyright @ 2006 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player